arXiv:1605.02070v1 [hep-ph] 6 May 2016

APFELgrid: a high performance tool for parton density determinations

Valerio Bertone[a], Stefano Carrazza[b], Nathan P. Hartland[a,*]

[a] *Rudolf Peierls Centre for Theoretical Physics, 1 Keble Road, University of Oxford, OX1 3NP, Oxford, UK*
[b] *Theoretical Physics Department, CERN, Geneva, Switzerland*

Abstract

We present a new software package designed to reduce the computational burden of hadron collider measurements in Parton Distribution Function (PDF) fits. The APFELgrid package converts interpolated weight tables provided by APPLgrid files into a more efficient format for PDF fitting by the combination with PDF and α_s evolution factors provided by APFEL. This combination significantly reduces the number of operations required to perform the calculation of hadronic observables in PDF fits and simplifies the structure of the calculation into a readily optimised scalar product. We demonstrate that our technique can lead to a substantial speed improvement when compared to existing methods without any reduction in numerical accuracy.

Keywords: QCD; parton distribution functions; fast predictions.

☆Preprints: OUTP-16-09P, CERN-TH-2016-103
*Corresponding author.
E-mail address: nathan.hartland@physics.ox.ac.uk

Program Summary

Manuscript Title: `APFELgrid`: a high performance tool for parton density determinations
Authors: V. Bertone, S. Carrazza, N.P. Hartland
Program Title: `APFELgrid`
Journal Reference:
Catalogue identifier:
Licensing provisions: MIT license
Programming language: C++
Computer: PC/Mac
Operating system: MacOS/Linux
RAM: varying
Keywords: QCD, PDF
Classification: 11.1 General, High Energy Physics and Computing
External routines/libraries: `APPLgrid`, `APFEL`
Nature of problem:
Fast computation of hadronic observables under the variation of parton distribution functions.
Solution method:
Combination of interpolated weight grids from `APPLgrid` files and evolution factors from `APFEL` into efficient `FastKernel` tables.
Running time: varying

1. Introduction

Measurements at colliders such as the Tevatron and Large Hadron Collider (LHC) have a unique capacity to shed light upon the internal dynamics of the proton and provide constraints upon proton PDFs [1]. However including large hadron collider datasets in PDF fits can provide a significant challenge due to the large computational footprint of performing accurate theoretical predictions over the many iterations required in a fitting procedure. In order to make the fullest use of current and future LHC results, efficient strategies for the computation of these observables must therefore be employed.

The typical Monte Carlo software packages used to perform predictions for hadron collider observables cannot be easily deployed in a PDF fit due the processing time required to obtain accurate results (usually of the order of a few hours or more per data point). To overcome such limitations, the typical strategy adopted for fast cross section prediction relies on the precomputation of the partonic hard cross sections in such a way that the standard numerical convolution with any set of PDFs can be reliably approximated by means of interpolation techniques.

Such interpolation strategies are implemented in the `APPLgrid` [2] and `FastNLO` [3] projects. For the computation of the hard cross sections, these packages rely on external codes to which they are interfaced by means of a suite of functions allowing for the filling of PDF- and α_s-independent look-up tables of cross section weights. Monte Carlo programs such as `MCFM` [4] and `NLOJet++` [5] have been interfaced directly to `APPLgrid`/`FastNLO` and more recently dedicated interfaces to automated general-purpose event generators have been developed. The `aMCfast` [6] and `MCgrid` [7] codes can generate interpolation grids in `APPLgrid`/`FastNLO` format by extracting the relevant information from the `MadGraph5_aMC@NLO` [8] and `SHERPA` [9] event generators respectively.

While these tools have proven to be invaluable in the extraction of parton densities, the volume of experimental data made available by LHC collaborations for use in PDF fits is already stretching the capabilities of the typical fitting technology. A standard global PDF fit may now include thousands of hadronic data points for which predictions have to be computed thousands of times during the minimisation process. As a consequence, performing these predictions using the standard interpolating tools, *i.e.* `APPLgrid` and `FastNLO`, starts to become prohibitively time-consuming. For this reason a high-performance tool tailored specifically to the requirements of PDF analysis becomes increasingly important.

The `FastKernel` method was developed to address this problem in the context of the NNPDF global analyses [10]. This method differs from the standard procedure à la `APPLgrid` or `FastNLO` in that it maximises the amount of information that is precomputed prior to fitting so as to minimise the amount of operations required during the fit. More specifically, the `FastKernel` method relies on the combination of precomputed hard cross sections with DGLAP evolution kernels into a single look-up table, here called a `FastKernel` (`FK`) table. In this way the prediction for a given hadronic observable can

be obtained by performing a simple matrix product between the respective FK table and PDFs evaluated directly at the fitting scale.

In this paper we present the APFELgrid package, a public implementation of the FastKernel method in which the hard partonic cross sections provided in an APPLgrid look-up table are combined with the DGLAP evolution kernels provided by the APFEL package [11].

This paper proceeds as follows. In Sect. 2 we present the technical details of the implementation of the FastKernel method. This is followed in Sect. 3 by a performance benchmark of the APFELgrid library and resulting FK tables. Finally, in Sect. 4 we summarise the results discussed in this work.

2. Interpolation tools for collider observables

Hadron collider observables are typically computed in QCD by means of a double convolution of parton densities with a hard scattering cross section. Consider for example the calculation of a general cross section $pp \to X$ with a set of PDFs $\{f\}$:

$$\sigma_{pp\to X} = \sum_s \sum_p \int dx_1\, dx_2\, \hat{\sigma}^{(p)(s)}\, \alpha_s^{p+p_{\rm LO}}(Q^2)\, F^{(s)}(x_1, x_2, Q^2)\,, \tag{1}$$

where Q^2 is the typical hard scale of the process, the index s sums over the active partonic subprocesses in the calculation, p sums over the perturbative orders used in the expansion, $p_{\rm LO}$ is the leading-order power of α_s for the process and $\hat{\sigma}^{(p)(s)}$ is the N^{p}LO contribution to the cross section for the partonic subprocess scattering $(s) \to X$. $F^{(s)}$ represents the subprocess parton density:

$$F^{(s)}(x_1, x_2, Q^2) = \sum_{i,j} C_{ij}^{(s)} f_i(x_1, Q^2) f_j(x_2, Q^2)\,, \tag{2}$$

where the $C_{ij}^{(s)}$ matrix enumerates the combinations of PDFs contributing to the s-th subprocess. The central observation of tools such as APPLgrid and FastNLO is that the PDF and α_S dependence may be factorised out of the convolution via expansion over a set of interpolating functions, spanning Q^2 and the two values of parton-x. For example one may represent the subprocess PDFs and α_S in terms of Lagrange basis polynomials $\mathcal{I}_\tau(Q^2)$, $\mathcal{I}_\alpha(x_1)$ and $\mathcal{I}_\beta(x_2)$ as:

$$\begin{aligned} &\alpha_s^{p+p_{\rm LO}}(Q^2)\, F^{(s)}(x_1, x_2, Q^2) = \\ &\sum_{\alpha,\beta,\tau} \alpha_s^{p+p_{\rm LO}}(Q_\tau^2)\, F_{\alpha\beta,\tau}^{(s)}\, \mathcal{I}_\tau(Q^2)\, \mathcal{I}_\alpha(x_1)\, \mathcal{I}_\beta(x_2), \end{aligned} \tag{3}$$

where we use the shorthand $F_{\alpha\beta,\tau}^{(s)} = F^{(s)}(x_\alpha, x_\beta, Q_\tau^2)$. Using these expressions in the double convolution of Eq. (1) one can finally obtain an expression for the desired cross section which depends upon the subprocess PDFs only through a simple product:

$$\sigma_{pp\to X} = \sum_p \sum_s \sum_{\alpha,\beta,\tau} \alpha_s^{p+p_{\rm LO}}(Q_\tau^2) W^{(p)(s)}_{\alpha\beta,\tau}\, F^{(s)}_{\alpha\beta,\tau}\,, \tag{4}$$

where

$$W^{(p)(s)}_{\alpha\beta,\tau} = \mathcal{I}_\tau(Q^2) \int dx_1\, dx_2\, \hat{\sigma}^{(p)(s)}\, \mathcal{I}_\alpha(x_1)\, \mathcal{I}_\beta(x_2)\,, \tag{5}$$

consists of the convolution of the hard cross section with the interpolating polynomials. This information may be stored in a precomputed look-up table. The final expression for the cross section in Eq. (4) is therefore a considerably simpler task to perform inside a fit than the direct evaluation of the double convolution.

2.1. The FastKernel method

A number of tools (*e.g.* `APFEL`, `HOPPET` [12] and `QCDNUM` [13]) are available which perform PDF evolution via an analogous interpolation procedure. In such a way PDFs at a general scale Q_τ may be expressed as a product of PDFs at some initial fitting scale Q_0 and an *evolution operator* obtained by the solution of the DGLAP equation.

$$f_i(x_\alpha, Q_\tau^2) = \sum_k \sum_\beta A^\tau_{\alpha\beta,ik}\; f_k(x_\beta, Q_0^2)\,, \tag{6}$$

where latin indices run over PDF flavour, greek indices run over points in an initial-scale interpolating x-grid and the evolution operator A may be accessed directly in the `APFEL` package. Given this operator, we may replace the (general-scale) PDFs used in the subprocess parton density Eq. (2) with their equivalent expressions evaluated at the fitting scale as

$$\begin{aligned} F^{(s)}_{\alpha\beta,\tau} &= \sum_{i,j} \sum_{k,l} \sum_{\delta,\gamma} C^{(s)}_{ij} \left[A^\tau_{\alpha\delta ik}\; f_k(x_\delta, Q_0^2) A^\tau_{\beta\gamma jl}\; f_l(x_\gamma, Q_0^2) \right] \\ &= \sum_{k,l} \sum_{\delta,\gamma} \widetilde{C}^{(s),\tau}_{kl,\alpha\beta\gamma\delta} f_k(x_\delta, Q_0^2) f_l(x_\gamma, Q_0^2)\,, \end{aligned} \tag{7}$$

with the object

$$\widetilde{C}^{(s),\tau}_{kl,\alpha\beta\gamma\delta} = \sum_{i,j} C^{(s)}_{ij}\, A^\tau_{\alpha\delta ik} A^\tau_{\beta\gamma jl}\,, \tag{8}$$

combining the operations of subprocess density construction and PDF evolution. Going further and using the expression for subprocess parton densities in Eq. (7) in the full cross section calculation we obtain

$$\sigma_{pp\to X} = \sum_{k,l} \sum_{\delta,\gamma} \sum_p \sum_s \sum_{\alpha,\beta} \sum_\tau \alpha_s^{p+p_{\rm LO}}(Q_\tau^2) W^{(p)(s)}_{\alpha\beta,\tau} \widetilde{C}^{(s),\tau}_{kl,\alpha\beta\gamma\delta} f_k(x_\delta, Q_0^2) f_l(x_\gamma, Q_0^2)\,. \tag{9}$$

Performing some further contractions it is possible to obtain an extremely compact expression for the calculation of the cross section in question, in terms of only the initial-scale PDFs and summing only over the initial scale interpolating x-grid and the incoming parton flavours:

$$\sigma_{pp\to X} = \sum_{k,l} \sum_{\delta,\gamma} \widetilde{W}_{kl,\delta\gamma}\, f_k(x_\delta, Q_0^2) f_l(x_\gamma, Q_0^2), \tag{10}$$

where the object:

$$\widetilde{W}_{kl,\delta\gamma} = \sum_{p} \sum_{s} \sum_{\alpha,\beta} \sum_{\tau} \alpha_s^{p+p_{\mathrm{LO}}}(Q_\tau^2) W^{(p)(s)}_{\alpha\beta,\tau} \widetilde{C}^{(s),\tau}_{kl,\alpha\beta\gamma\delta} \tag{11}$$

is referred to here as an `FK` table, and combines the information stored in `APPLgrid`-style interpolated weight grids with analogously interpolated DGLAP evolution operators. This combination enables for a maximally efficient expression for the calculation of observables at hadron colliders under PDF variation, without invoking any additional approximation.

2.2. Features and limitations of `FK` *tables*

The `FK` product of Eq. (10) differs with respect to the product in Eq. (4) in several notable ways. Firstly the typical `APPLgrid` or `FastNLO` products use as input PDFs at a general scale, requiring that PDF evolution *e.g.* Eq. (6) be performed for every variation of the PDFs during the fit. In the `FK` product this evolution is pre-cached at the stage of `FK` table generation, requiring only initial-scale PDFs at the time of fitting. This pre-caching of the evolution also removes the need to sum over hard scale and perturbative order during the fit, further reducing the number of operations required. As the `FK` product acts directly at the fitting scale, it benefits from the typically reduced number of active partonic modes, with the sum over flavours in Eq. (10) being limited to those directly parametrised in the fit. Having reduced the calculation to such a simple form, it is also straightforward to apply standard computational tools such as multi-threading through e.g OpenMP or Single Instruction Multiple Data (SIMD) operations such as SSE or AVX to further reduce computational expense.

While these features provide significant performance enhancements, the `FK` table format is not suitable as a complete replacement for tools such as `APPLgrid`. The pre-computation of the PDF evolution necessarily means that all theory parameters such as perturbative order, strong coupling and factorization/renormalization scales are inextricably embedded in each `FK` table. In order to perform PDF fits including variations of these parameters, multiple `FK` tables must be computed, each with different theory settings. While performing such a re-calculation directly from Monte Carlo codes would be exceptionally time consuming, the data representation in `APPLgrid` files allows for an efficient (re)-combination with varying theory parameters.

3. Performance benchmarks

We shall now examine the performance differences between the two expressions for fast interpolated cross section prediction Eq. (4) (`APPLgrid`) and Eq. (10) (`FK`). In order to provide a comprehensive benchmark, we consider here a wide range of processes including LHC and Tevatron electroweak vector boson production measurements [14–22], $t\bar{t}$ total cross sections [23–26], double-differential Drell-Yan cross sections [27, 28] and inclusive jet data [29–32]. Predictions are performed over a wide range of kinematics, for a total of 52 source `APPLgrid` files corresponding to the majority of available LHC and Tevatron datasets applicable to PDF determination. While the source `APPLgrid` files have varying grid densities in x and Q^2, for the purposes of comparison the corresponding `FK` tables are produced consistently with 30 points in x, and at an initial scale below the charm threshold, therefore with seven active partonic species. These settings are chosen so as to provide a realistic comparison, in a production environment the density and distribution of the x-grid may be adjusted to match interpolation accuracy requirements. For these comparisons the `FK` table is stored as double-precision in memory for table generation and in single-precision for the purposes of computing the `FK` product.

In Fig. 1 we compare the average time taken per datapoint for the `FK` and `APPLgrid` calculations, for all of the 52 tables. We show timings for the `FK` calculation in four different configurations: AVX-OpenMP 2x (2 CPU cores), AVX, SSE3 and the standard double precision product. Due to the inherent structural differences between the `FK` and `APPLgrid` procedures, results from the `FK` calculation are systematically faster than those from `APPLgrid`. In particular, when comparing `FK` AVX-OpenMP 2x to `APPLgrid` timings we obtain minimally a factor of ten improvement in speed for electroweak vector boson production and a maximum factor of 2000 improvement in predictions for inclusive jet data. Across all processes and kinematic regions we observe significant performance improvements from using the `FK` calculation even without the use of SIMD or multi-threading.

While sheer computational speed is typically the primary consideration when computing observables in a PDF fit, other factors such as table size in the filesystem and memory, along with the computational cost of pre-computing `FK` tables must be considered. Indeed, the computation of the `FK` table in Eq. (11) requires a great deal of operations which can be time consuming, particularly in the case of source `APPLgrid` files with very high interpolation precision.

In Fig. 2 we examine the `FK` table generation time with `APFELgrid`, `FK` table file size and memory usage of the `FK` tables arising from the same source `APPLgrid` files as discussed in Fig. 1. When examining the table generation time per point, we observe timings from a few milliseconds to 3.5 minutes per point, with differences arising from the varying grid densities used in the input `APPLgrid` files. In terms of the grid size on disk, `FK` tables are typically larger than their corresponding `APPLgrid` files, primarily as the `FK` file format is encoded in plain text for compatibility whereas `APPLgrid` files are expressed in binary

as ROOT files. However when measuring the in-memory resident set size used by the two procedures, the amount of system memory used by FK tables is systematically less than APPLgrid files for all processes considered here by at least 75%. Note that this effect is in part due to the differing precisions of the default representations.

4. Conclusion

In this work we have demonstrated that by employing the so-called FastKernel method, it is possible to convert an APPLgrid weight table into a derived format, referred to as an FK table, including the effects of PDF and α_s evolution. This procedure has been implemented in the APFELgrid package, supplied as a set of C++ routines designed to supplement the PDF evolution library APFEL with FK table generation capabilities. The APFELgrid package allows one to obtain a computationally efficient expression for the calculation of hadronic cross sections, in terms of only the initial-scale PDFs and summing only over the initial scale interpolating x-grid and the incoming parton flavours. The simple structure of the resulting product makes FK tables particularly suitable for the efficient use of computational tools such as SIMD and OpenMP.

We have shown that in several practical examples the numerical evaluation of an FK product is considerably faster than the corresponding APPLgrid product, even in the case where neither SIMD or multi-threading are applied. FK tables are supplied in a simple plain-text format in order to simplify the construction of user interfaces, and therefore are typically larger than corresponding APPLgrids. However we have shown that the in-memory resident set sizes occupied by FK tables are typically smaller than those required by APPLgrids, in our examples by at least 75%.

The substantial speed improvement of FK tables with respect to APPLgrid in association with the reduction in memory footprint makes the APFELgrid code a valuable tool for modern PDF fits including large collider datasets.

The APFELgrid package and associated documentation are publicly available from the webpage:

https://github.com/nhartland/APFELgrid

Acknowledgements

The authors would like to thank members of the NNPDF Collaboration for their support and motivation for this work; particularly Juan Rojo, Luigi del Debbio and Alberto Guffanti. We would also like to thank Mark Sutton for helpful comments on the paper. V. B. and N. H. are supported by an European Research Council Starting Grant "PDF4BSM". S. C. is supported by the HICCUP ERC Consolidator grant (614577).

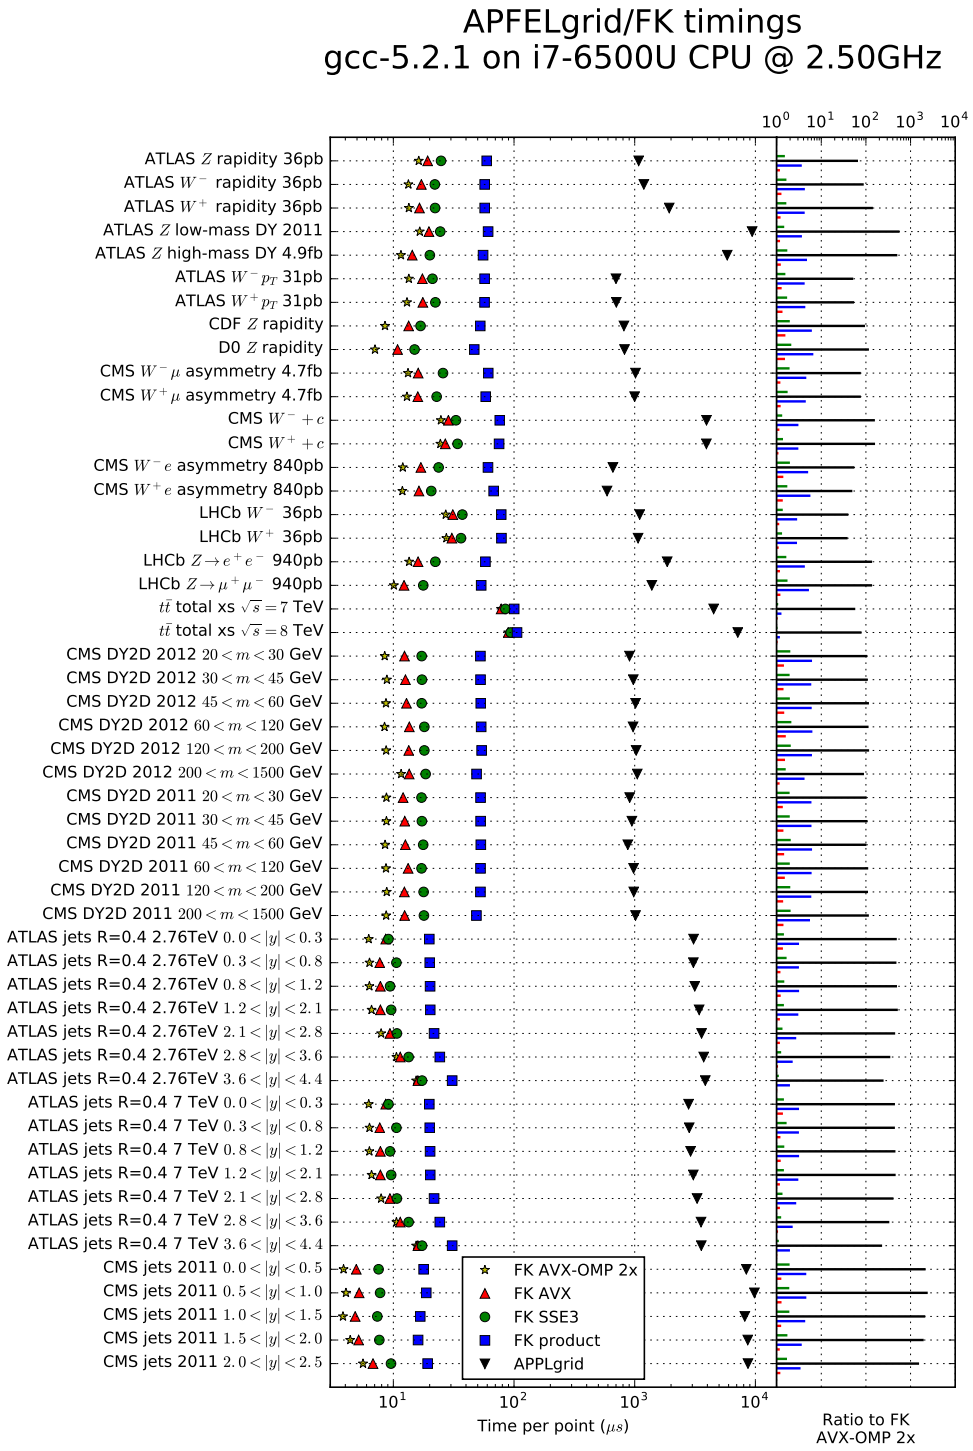
APFELgrid/FK timings
gcc-5.2.1 on i7-6500U CPU @ 2.50GHz
ATLAS Z rapidity 36pb
ATLAS W^- rapidity 36pb
ATLAS W^+ rapidity 36pb
ATLAS Z low-mass DY 2011
ATLAS Z high-mass DY 4.9fb
ATLAS $W^- p_T$ 31pb
ATLAS $W^+ p_T$ 31pb
CDF Z rapidity
D0 Z rapidity
CMS $W^- \mu$ asymmetry 4.7fb
CMS $W^+ \mu$ asymmetry 4.7fb
CMS $W^- + c$
CMS $W^+ + c$
CMS $W^- e$ asymmetry 840pb
CMS $W^+ e$ asymmetry 840pb
LHCb W^- 36pb
LHCb W^+ 36pb
LHCb $Z \to e^+ e^-$ 940pb
LHCb $Z \to \mu^+ \mu^-$ 940pb
$t\bar{t}$ total xs $\sqrt{s} = 7$ TeV
$t\bar{t}$ total xs $\sqrt{s} = 8$ TeV
CMS DY2D 2012 $20 < m < 30$ GeV
CMS DY2D 2012 $30 < m < 45$ GeV
CMS DY2D 2012 $45 < m < 60$ GeV
CMS DY2D 2012 $60 < m < 120$ GeV
CMS DY2D 2012 $120 < m < 200$ GeV
CMS DY2D 2012 $200 < m < 1500$ GeV
CMS DY2D 2011 $20 < m < 30$ GeV
CMS DY2D 2011 $30 < m < 45$ GeV
CMS DY2D 2011 $45 < m < 60$ GeV
CMS DY2D 2011 $60 < m < 120$ GeV
CMS DY2D 2011 $120 < m < 200$ GeV
CMS DY2D 2011 $200 < m < 1500$ GeV
ATLAS jets R=0.4 2.76TeV $0.0 < |y| < 0.3$
ATLAS jets R=0.4 2.76TeV $0.3 < |y| < 0.8$
ATLAS jets R=0.4 2.76TeV $0.8 < |y| < 1.2$
ATLAS jets R=0.4 2.76TeV $1.2 < |y| < 2.1$
ATLAS jets R=0.4 2.76TeV $2.1 < |y| < 2.8$
ATLAS jets R=0.4 2.76TeV $2.8 < |y| < 3.6$
ATLAS jets R=0.4 2.76TeV $3.6 < |y| < 4.4$
ATLAS jets R=0.4 7 TeV $0.0 < |y| < 0.3$
ATLAS jets R=0.4 7 TeV $0.3 < |y| < 0.8$
ATLAS jets R=0.4 7 TeV $0.8 < |y| < 1.2$
ATLAS jets R=0.4 7 TeV $1.2 < |y| < 2.1$
ATLAS jets R=0.4 7 TeV $2.1 < |y| < 2.8$
ATLAS jets R=0.4 7 TeV $2.8 < |y| < 3.6$
ATLAS jets R=0.4 7 TeV $3.6 < |y| < 4.4$
CMS jets 2011 $0.0 < |y| < 0.5$
CMS jets 2011 $0.5 < |y| < 1.0$
CMS jets 2011 $1.0 < |y| < 1.5$
CMS jets 2011 $1.5 < |y| < 2.0$
CMS jets 2011 $2.0 < |y| < 2.5$
10^0 10^1 10^2 10^3 10^4
FK AVX-OMP 2x
FK AVX
FK SSE3
FK product
APPLgrid
10^1 10^2 10^3 10^4
Time per point (μs)
Ratio to FK AVX-OMP 2x

Figure 1: Performance comparisons between `FK` with AVX-OpenMP 2x, AVX, SSE3, double precision convolution and `APPLgrid` convolution time per point and process.

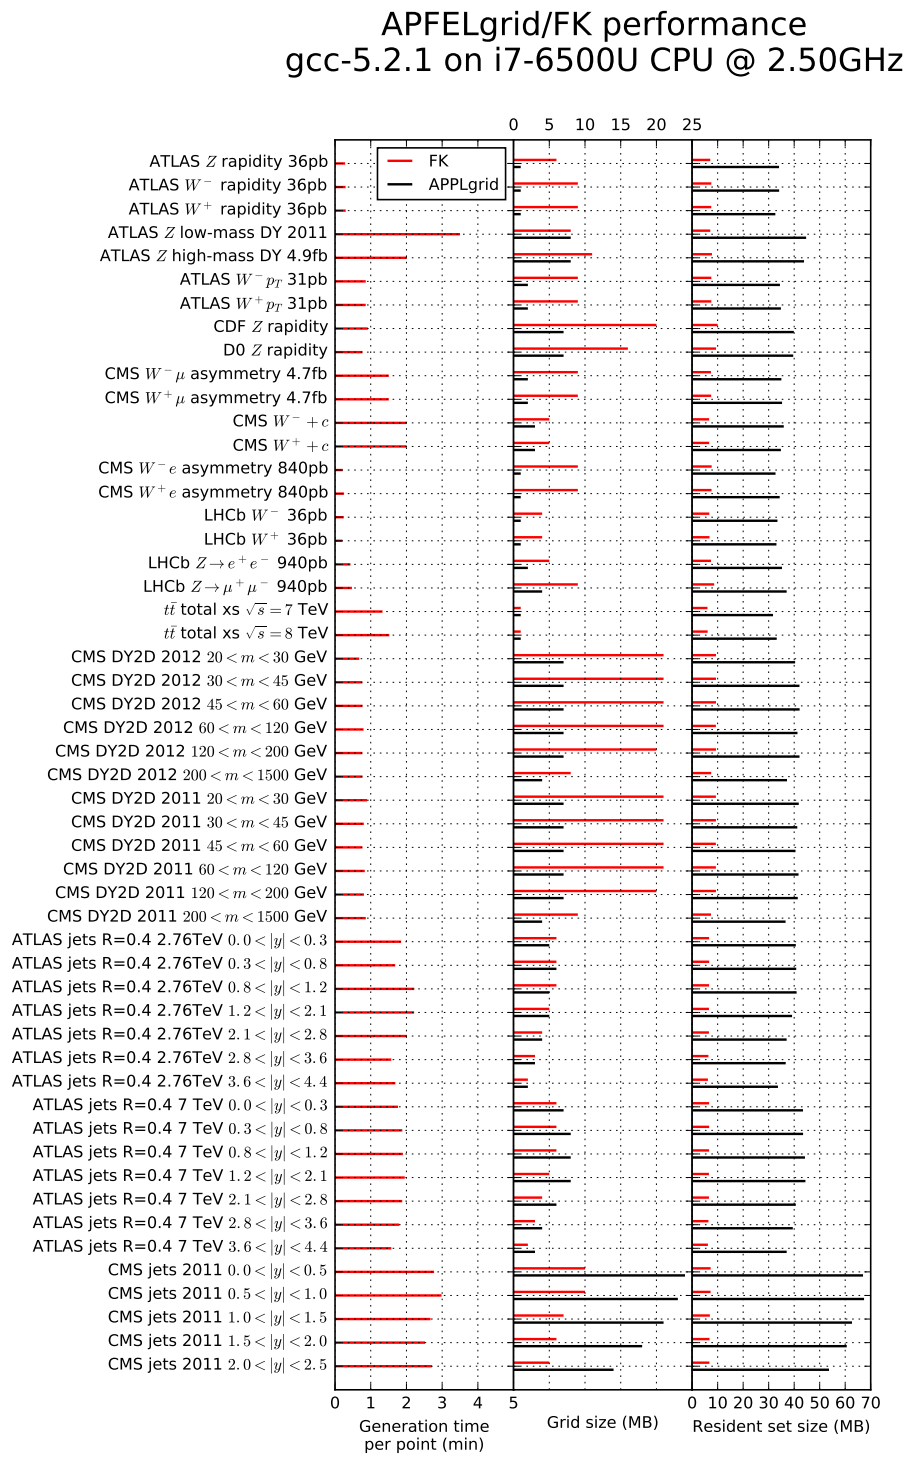
APFELgrid/FK performance
gcc-5.2.1 on i7-6500U CPU @ 2.50GHz
FK
APPLgrid
ATLAS Z rapidity 36pb
ATLAS W^- rapidity 36pb
ATLAS W^+ rapidity 36pb
ATLAS Z low-mass DY 2011
ATLAS Z high-mass DY 4.9fb
ATLAS $W^- p_T$ 31pb
ATLAS $W^+ p_T$ 31pb
CDF Z rapidity
D0 Z rapidity
CMS $W^- \mu$ asymmetry 4.7fb
CMS $W^+ \mu$ asymmetry 4.7fb
CMS $W^- + c$
CMS $W^+ + c$
CMS $W^- e$ asymmetry 840pb
CMS $W^+ e$ asymmetry 840pb
LHCb W^- 36pb
LHCb W^+ 36pb
LHCb $Z \to e^+ e^-$ 940pb
LHCb $Z \to \mu^+ \mu^-$ 940pb
$t\bar{t}$ total xs $\sqrt{s} = 7$ TeV
$t\bar{t}$ total xs $\sqrt{s} = 8$ TeV
CMS DY2D 2012 $20 < m < 30$ GeV
CMS DY2D 2012 $30 < m < 45$ GeV
CMS DY2D 2012 $45 < m < 60$ GeV
CMS DY2D 2012 $60 < m < 120$ GeV
CMS DY2D 2012 $120 < m < 200$ GeV
CMS DY2D 2012 $200 < m < 1500$ GeV
CMS DY2D 2011 $20 < m < 30$ GeV
CMS DY2D 2011 $30 < m < 45$ GeV
CMS DY2D 2011 $45 < m < 60$ GeV
CMS DY2D 2011 $60 < m < 120$ GeV
CMS DY2D 2011 $120 < m < 200$ GeV
CMS DY2D 2011 $200 < m < 1500$ GeV
ATLAS jets R=0.4 2.76TeV $0.0 < |y| < 0.3$
ATLAS jets R=0.4 2.76TeV $0.3 < |y| < 0.8$
ATLAS jets R=0.4 2.76TeV $0.8 < |y| < 1.2$
ATLAS jets R=0.4 2.76TeV $1.2 < |y| < 2.1$
ATLAS jets R=0.4 2.76TeV $2.1 < |y| < 2.8$
ATLAS jets R=0.4 2.76TeV $2.8 < |y| < 3.6$
ATLAS jets R=0.4 2.76TeV $3.6 < |y| < 4.4$
ATLAS jets R=0.4 7 TeV $0.0 < |y| < 0.3$
ATLAS jets R=0.4 7 TeV $0.3 < |y| < 0.8$
ATLAS jets R=0.4 7 TeV $0.8 < |y| < 1.2$
ATLAS jets R=0.4 7 TeV $1.2 < |y| < 2.1$
ATLAS jets R=0.4 7 TeV $2.1 < |y| < 2.8$
ATLAS jets R=0.4 7 TeV $2.8 < |y| < 3.6$
ATLAS jets R=0.4 7 TeV $3.6 < |y| < 4.4$
CMS jets 2011 $0.0 < |y| < 0.5$
CMS jets 2011 $0.5 < |y| < 1.0$
CMS jets 2011 $1.0 < |y| < 1.5$
CMS jets 2011 $1.5 < |y| < 2.0$
CMS jets 2011 $2.0 < |y| < 2.5$
0 5 10 15 20 25
0 1 2 3 4
5
0 10 20 30 40 50 60 70
Generation time per point (min)
Grid size (MB)
Resident set size (MB)

Figure 2: `APFELgrid` generation time per point and comparison to `APPLgrid` for grid size on disk and resident set size (RSS).

References

[1] J. Rojo, et al., The PDF4LHC report on PDFs and LHC data: Results from Run I and preparation for Run II, J. Phys. G42 (2015) 103103. arXiv:1507.00556, doi:10.1088/0954-3899/42/10/103103.

[2] T. Carli, et al., A posteriori inclusion of parton density functions in NLO QCD final-state calculations at hadron colliders: The APPLGRID Project, Eur.Phys.J. C66 (2010) 503. arXiv:0911.2985, doi:10.1140/epjc/s10052-010-1255-0.

[3] M. Wobisch, D. Britzger, T. Kluge, K. Rabbertz, F. Stober, Theory-Data Comparisons for Jet Measurements in Hadron-Induced ProcessesarXiv:1109.1310.

[4] J. M. Campbell, R. K. Ellis, MCFM for the Tevatron and the LHC, Nucl. Phys. Proc. Suppl. 205-206 (2010) 10–15. arXiv:1007.3492, doi:10.1016/j.nuclphysbps.2010.08.011.

[5] Z. Nagy, Next-to-leading order calculation of three jet observables in hadron hadron collision, Phys. Rev. D68 (2003) 094002. arXiv:hep-ph/0307268, doi:10.1103/PhysRevD.68.094002.

[6] V. Bertone, R. Frederix, S. Frixione, J. Rojo, M. Sutton, aMCfast: automation of fast NLO computations for PDF fits, JHEP 08 (2014) 166. arXiv:1406.7693, doi:10.1007/JHEP08(2014)166.

[7] L. Del Debbio, N. P. Hartland, S. Schumann, MCgrid: projecting cross section calculations on grids, Comput. Phys. Commun. 185 (2014) 2115–2126. arXiv:1312.4460, doi:10.1016/j.cpc.2014.03.023.

[8] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, JHEP 07 (2014) 079. arXiv:1405.0301, doi:10.1007/JHEP07(2014)079.

[9] T. Gleisberg, S. Hoeche, F. Krauss, M. Schonherr, S. Schumann, F. Siegert, J. Winter, Event generation with SHERPA 1.1, JHEP 02 (2009) 007. arXiv:0811.4622, doi:10.1088/1126-6708/2009/02/007.

[10] R. D. Ball, et al., Parton distributions for the LHC Run II, JHEP 04 (2015) 040. arXiv:1410.8849, doi:10.1007/JHEP04(2015)040.

[11] V. Bertone, S. Carrazza, J. Rojo, APFEL: A PDF Evolution Library with QED corrections, Comput. Phys. Commun. 185 (2014) 1647–1668. arXiv:1310.1394, doi:10.1016/j.cpc.2014.03.007.

[12] G. P. Salam, J. Rojo, A Higher Order Perturbative Parton Evolution Toolkit (HOPPET), Comput. Phys. Commun. 180 (2009) 120–156. arXiv:0804.3755, doi:10.1016/j.cpc.2008.08.010.

[13] M. Botje, QCDNUM: Fast QCD Evolution and Convolution, Comput. Phys. Commun. 182 (2011) 490–532. arXiv:1005.1481, doi:10.1016/j.cpc.2010.10.020.

[14] R. Aaij, et al., Measurement of the cross-section for $Z \to e^+e^-$ production in pp collisions at $\sqrt{s} = 7$ TeV, JHEP 02 (2013) 106. arXiv:1212.4620, doi:10.1007/JHEP02(2013)106.

[15] R. Aaij, et al., Inclusive W and Z production in the forward region at $\sqrt{s} = 7$ TeV, JHEP 06 (2012) 058. arXiv:1204.1620, doi:10.1007/JHEP06(2012)058.

[16] S. Chatrchyan, et al., Measurement of the muon charge asymmetry in inclusive $pp \to W + X$ production at $\sqrt{s} = 7$ TeV and an improved determination of light parton distribution functions, Phys. Rev. D90 (3) (2014) 032004. arXiv:1312.6283, doi:10.1103/PhysRevD.90.032004.

[17] S. Chatrchyan, et al., Measurement of associated W + charm production in pp collisions at $\sqrt{s} = 7$ TeV, JHEP 02 (2014) 013. arXiv:1310.1138, doi:10.1007/JHEP02(2014)013.

[18] S. Chatrchyan, et al., Measurement of the electron charge asymmetry in inclusive W production in pp collisions at $\sqrt{s} = 7$ TeV, Phys. Rev. Lett. 109 (2012) 111806. arXiv:1206.2598, doi:10.1103/PhysRevLett.109.111806.

[19] G. Aad, et al., Measurement of the high-mass Drell–Yan differential cross-section in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector, Phys. Lett. B725 (2013) 223–242. arXiv:1305.4192, doi:10.1016/j.physletb.2013.07.049.

[20] G. Aad, et al., Measurement of the Transverse Momentum Distribution of W Bosons in pp Collisions at $\sqrt{s} = 7$ TeV with the ATLAS Detector, Phys. Rev. D85 (2012) 012005. arXiv:1108.6308, doi:10.1103/PhysRevD.85.012005.

[21] G. Aad, et al., Measurement of the inclusive $W^\pm$ and Z/γ cross sections in the electron and muon decay channels in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector, Phys. Rev. D85 (2012) 072004. arXiv:1109.5141, doi:10.1103/PhysRevD.85.072004.

[22] T. A. Aaltonen, et al., Measurement of $d\sigma/dy$ of Drell-Yan e^+e^- pairs in the Z Mass Region from $p\bar{p}$ Collisions at $\sqrt{s} = 1.96$ TeV, Phys. Lett. B692 (2010) 232–239. arXiv:0908.3914, doi:10.1016/j.physletb.2010.06.043.

[23] G. Aad, et al., Measurement of the cross section for top-quark pair production in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector using final states with two high-pt leptons, JHEP 05 (2012) 059. arXiv:1202.4892, doi:10.1007/JHEP05(2012)059.

[24] S. Chatrchyan, et al., Measurement of the $t\bar{t}$ production cross section in the dilepton channel in pp collisions at $\sqrt{s} = 8$ TeV, JHEP 02 (2014) 024, [Erratum: JHEP02,102(2014)]. arXiv:1312.7582, doi:10.1007/JHEP02(2014)024,10.1007/JHEP02(2014)102.

[25] S. Chatrchyan, et al., Measurement of the $t\bar{t}$ production cross section in the dilepton channel in pp collisions at $\sqrt{s} = 7$ TeV, JHEP 11 (2012) 067. arXiv:1208.2671, doi:10.1007/JHEP11(2012)067.

[26] S. Chatrchyan, et al., Measurement of the $t\bar{t}$ production cross section in pp collisions at $\sqrt{s} = 7$ TeV with lepton + jets final states, Phys. Lett. B720 (2013) 83–104. arXiv:1212.6682, doi:10.1016/j.physletb.2013.02.021.

[27] S. Chatrchyan, et al., Measurement of the differential and double-differential Drell-Yan cross sections in proton-proton collisions at $\sqrt{s} = 7$ TeV, JHEP 12 (2013) 030. arXiv:1310.7291, doi:10.1007/JHEP12(2013)030.

[28] V. Khachatryan, et al., Measurements of differential and double-differential Drell-Yan cross sections in proton-proton collisions at 8 TeV, Eur. Phys. J. C75 (4) (2015) 147. arXiv:1412.1115, doi:10.1140/epjc/s10052-015-3364-2.

[29] S. Chatrchyan, et al., Measurements of differential jet cross sections in proton-proton collisions at $\sqrt{s} = 7$ TeV with the CMS detector, Phys. Rev. D87 (11) (2013) 112002, [Erratum: Phys. Rev.D87,no.11,119902(2013)]. arXiv:1212.6660, doi:10.1103/PhysRevD.87.112002,10.1103/PhysRevD.87.119902.

[30] G. Aad, et al., Measurement of inclusive jet and dijet production in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector, Phys. Rev. D86 (2012) 014022. arXiv:1112.6297, doi:10.1103/PhysRevD.86.014022.

[31] G. Aad, et al., Measurement of the inclusive jet cross section in pp collisions at $\sqrt{s} = 2.76$ TeV and comparison to the inclusive jet cross section at $\sqrt{s} = 7$ TeV using the ATLAS detector, Eur. Phys. J. C73 (8) (2013) 2509. arXiv:1304.4739, doi:10.1140/epjc/s10052-013-2509-4.

[32] V. M. Abazov, et al., Measurement of the shape of the boson rapidity distribution for $p\bar{p} \to Z/\gamma^* \to e^+e^- + X$ events produced at $\sqrt{s}$ of 1.96-TeV, Phys. Rev. D76 (2007) 012003. arXiv:hep-ex/0702025, doi:10.1103/PhysRevD.76.012003.